imply with UserEvidence

THE BREAKING POINT FOR

Observability leaders.

Why rising cost and scale limits are driving a new generation of observability architectures

> **“We’re constantly having to choose between controlling costs and keeping the data we need. Splunk just doesn’t scale with the volume we’re dealing with.”**
>
> Head of Cloud and Security Operations – Global Dental Technology Company

Executive summary

Observability—once the cornerstone of proactive monitoring and troubleshooting—is straining under the weight of modern data growth. The scale and complexity of today’s event data are overwhelming architectures built for a smaller, simpler era.

In our survey of senior observability leaders across large enterprises, one pattern was clear: organizations are spending more, retaining less, and making constant trade-offs between visibility, cost, and agility.

01. Spending is up, satisfaction is down.

While 62% now spend more than a quarter of their observability budget on Splunk, only 13% are “very satisfied” with Splunk.

02. Tradeoffs are unavoidable.

Nearly 80% are filtering, archiving, or offloading logs to control costs—moves that risk degraded query performance and blind spots in real-time threat detection.

03. Leaders are ready to move

87% are actively looking or open to Splunk-compatible alternatives that ease cost and scale pressures.

The path forward is a new model—a unified, decoupled data layer purpose-built for today’s observability realities. Imply Lumi, the industry’s first Observability Warehouse, delivers scale, speed, and savings while allowing teams to keep the tools and workflows they already trust.

Introduction

For more than a decade, observability has been central to how organizations monitor complex systems, interpret event data, and act proactively. But today's data volumes are testing the limits of the architectures that made observability possible.

Many legacy platforms were designed in a different era—before the explosion of distributed systems, cloud workloads, and AI-driven analytics. As a result, they struggle to adapt to modern scale, complexity, and cost efficiency.

To understand how organizations are responding, we partnered with independent research firm UserEvidence to survey senior observability and platform leaders from large enterprises—most with more than 10,000 employees.

The findings are clear: teams are spending more but realizing less value. Rising data volumes, escalating costs, and architectural bottlenecks are creating tough trade-offs between visibility, performance, and affordability.

This report examines those trade-offs and explores how a modern, decoupled data architecture—exemplified by Imply Lumi, the industry's first Observability Warehouse—offers a new path forward. Lumi helps organizations escape the cycle of rising costs and shrinking visibility by delivering scale, speed, and savings without changing the tools or workflows teams rely on.

Key findings:
Overspending and decreasing satisfaction

Overspending is commonplace

Splunk is incredibly costly to the mature, enterprise organizations we surveyed. Every respondent reported spending more than **$100k/year** on Splunk, and **35% spend more than $1M.**

These numbers represent a huge portion of any observability budget, even at large companies: **62% say more than a quarter of their observability and security budget goes to Splunk costs.**

Roughly what portion of your observability or security budget is allocated to Splunk?

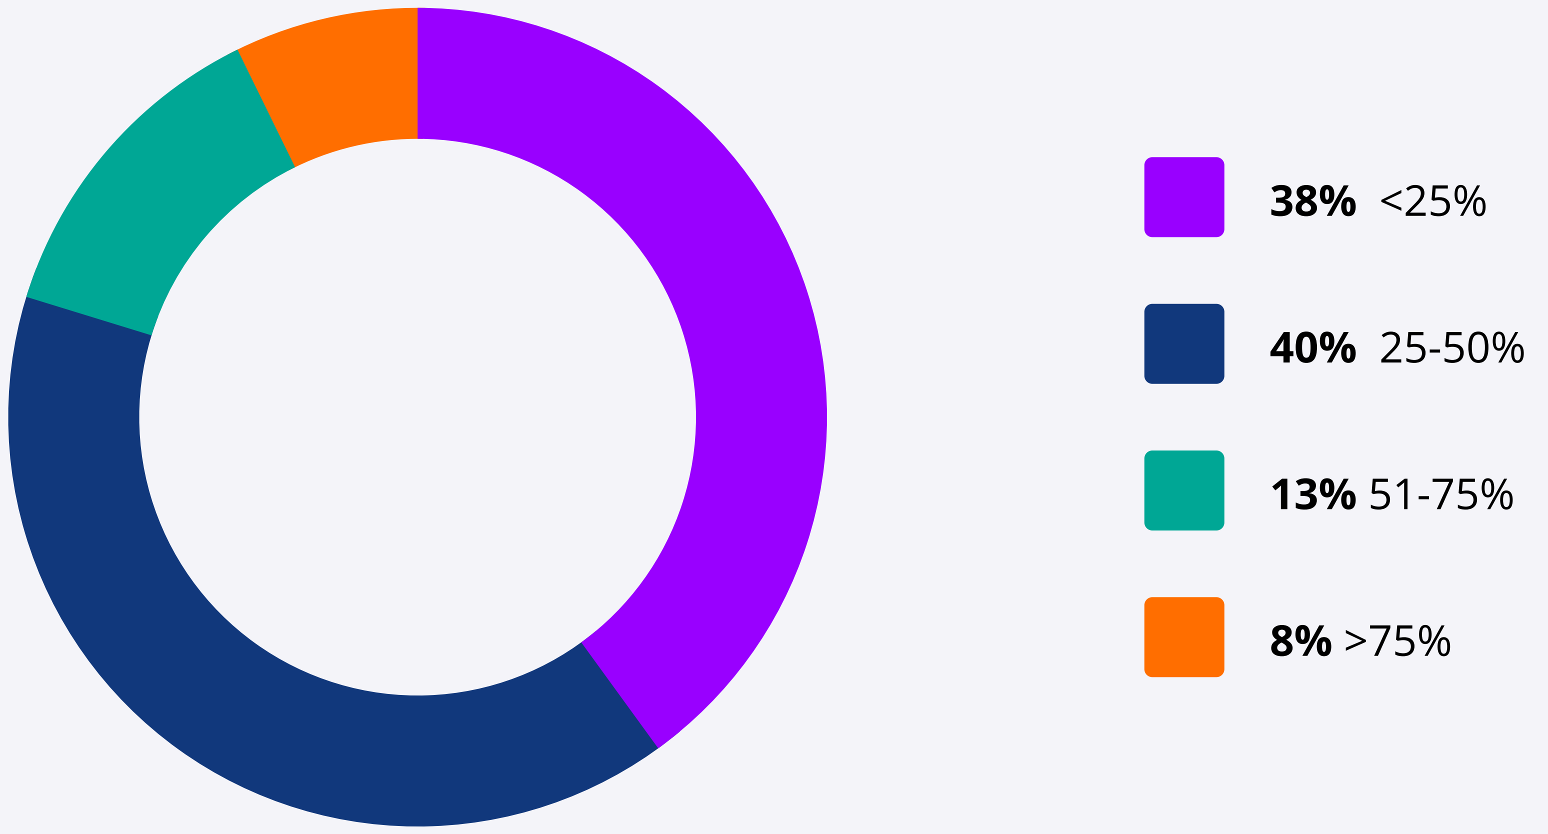

Executives consistently highlighted the strain this creates:

> **“As we put more and more data [into Splunk], the cost just skyrockets and we have a lot of budget constraints right now. So that makes it even more crucial that we manage our budgets more efficiently.”**

Director of Data Engineering and AI/ML Architecture – Large Public University System

> **“Splunk’s licensing model is a tax on growth. Every new use case, every additional workload, just means higher spend.”**

Kyle Nosker | Associate Director, Sourcing Operations Management, Merck

Rising data, rising costs

Overspending is on course to continue with increasing data usage. **65% of respondents increased their Splunk usage in the past year,** which in turn drove up both hardware and licensing costs.

How has your Splunk usage changed over the last 12 months?

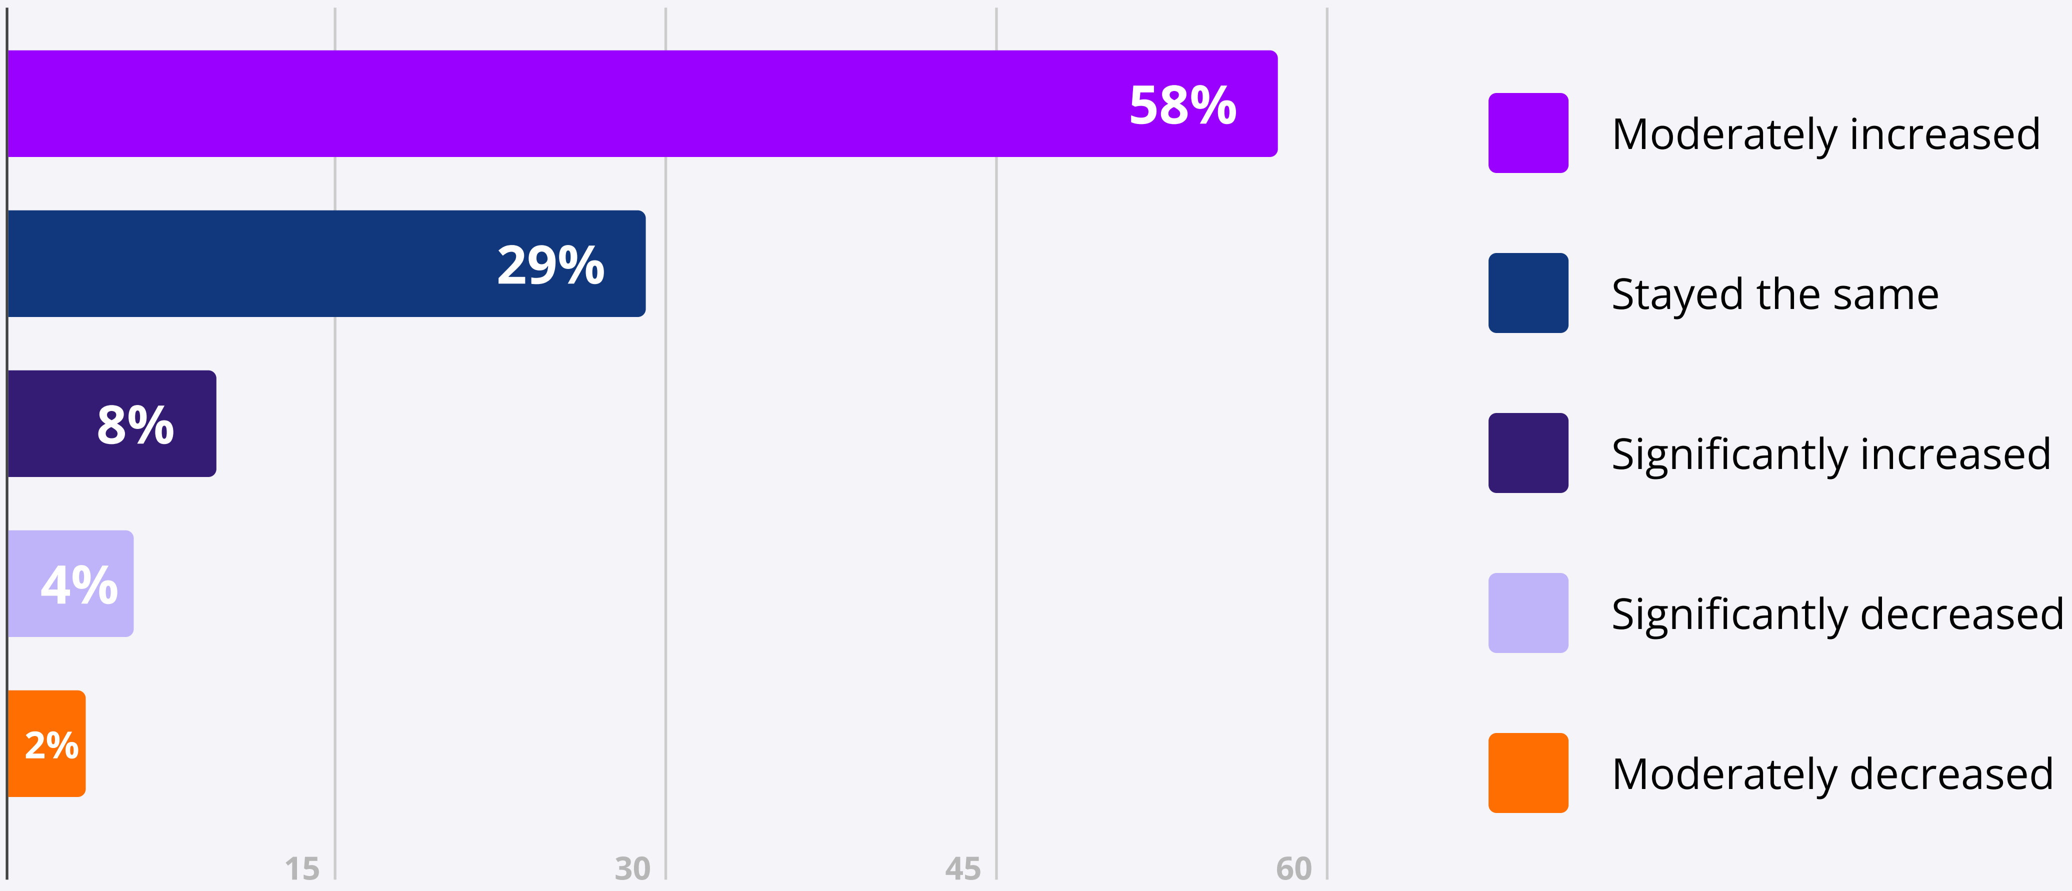

And there’s no sign of data slowing down. In an independent study of observability practitioners, **100% of respondents reported significant data growth in the past three years,** with two-thirds saying data ingestion volumes had tripled in that period. For Splunk customers, more data always equals more cost.

> **“We’re constantly having to choose between controlling costs and keeping the data we need. Splunk just doesn’t scale with the volume we’re dealing with.”**

Eike Schuster | Head of Cloud & Security Operations, Dentsply Sirona

Retention limits hurt visibility

To keep Splunk costs under control, most orgs are forced to reduce the amount of data they retain. **Nearly 80% of executives say they are filtering, archiving, or offloading logs** just to keep their Splunk bill manageable.

Have you implemented any strategies to reduce your Splunk Bill?

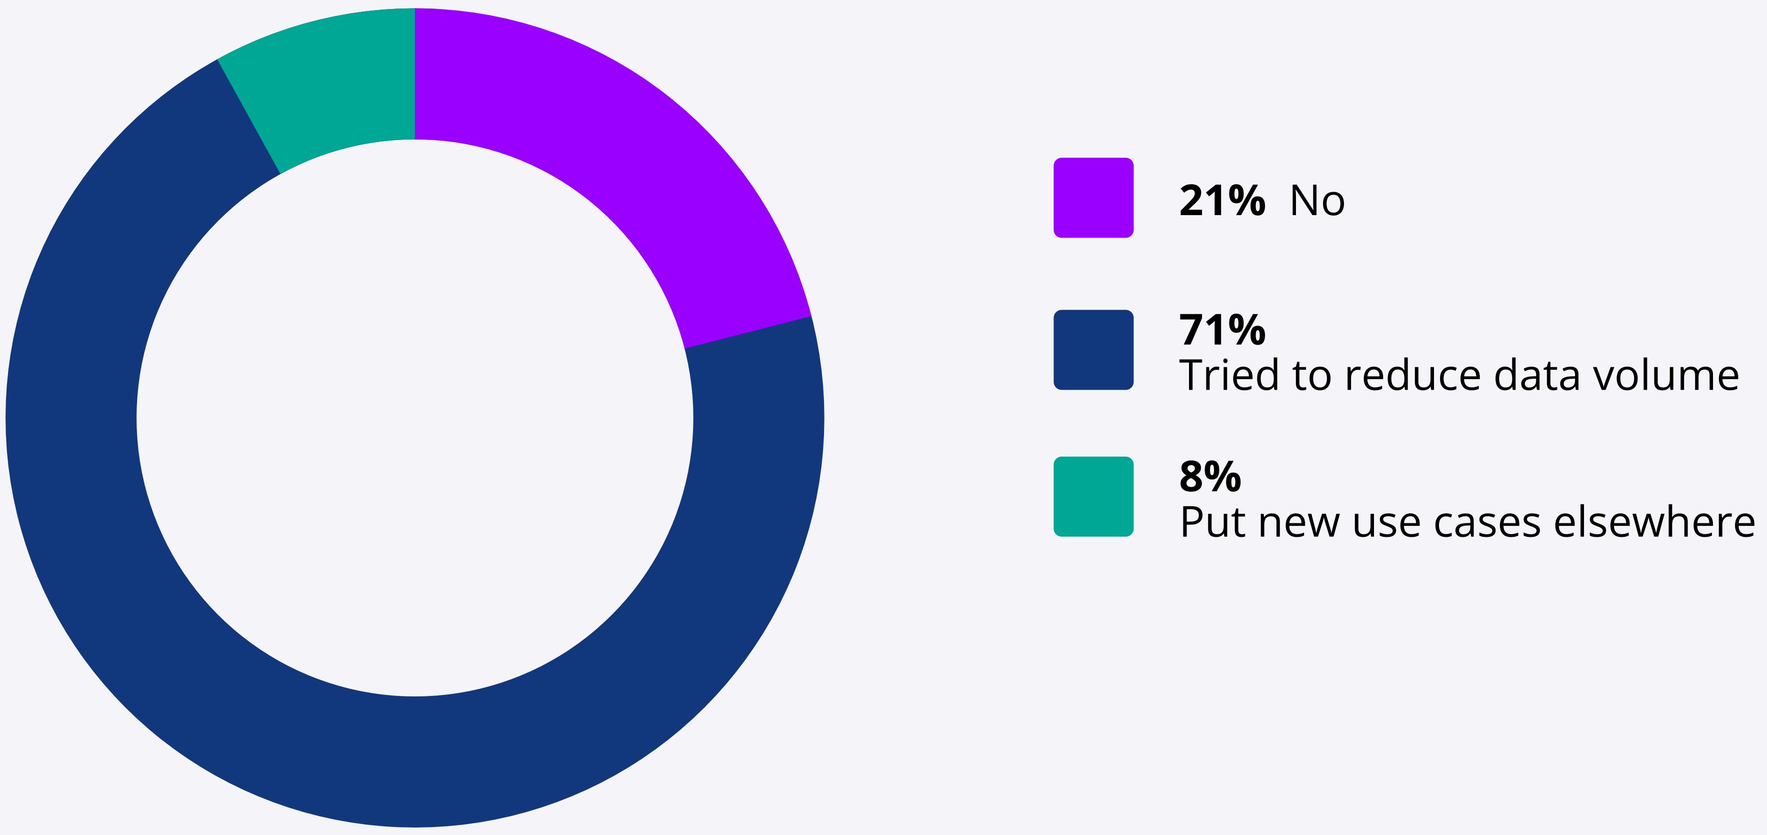

Tradeoffs are painful. Rehydrating archived data is slow and efficient, leaving teams unable to respond quickly during critical investigations. Meanwhile, filtering data before ingestion risks missing valuable signals altogether.

> **“We constantly hit retention walls. Historical investigations are painful because the data just isn’t there, or it’s locked away.”**
>
> **Kyle Nosker** | Merck

Survey data backs this up: 87% of respondents say slow query performance, caused by inaccessible data, disrupts critical use cases like threat detection, incident response, and executive reporting.

Practitioners described the impact firsthand:

> **“Rehydrating archived Splunk data is painful. It takes so much time that by the time you get access to the logs, the investigation has already moved on.”**

Tony Hashem | CISO & Director of IT Security, Vermont State Colleges System

> **“We push more data to cold storage than we’d like, and we filter out logs that might be valuable later — just to keep costs in check. Honestly, it’s not the way I’d like to run things. But when you're trying to balance limited resources and rising storage costs, you're often forced to settle for less insight than you'd want.”**

Jon Scarpa | Director, Information Technology at RDAbbot

Satisfaction is eroding and execs are looking elsewhere

As costs climb and retention shrinks, executives are increasingly dissatisfied with the value they're getting from Splunk. **Only 13% of respondents described themselves as "very satisfied" with the cost-to-value ratio.** The majority—three-quarters—say they are merely neutral or somewhat satisfied, underscoring growing frustration.

Given that Splunk often consumes more than a quarter of the observability budget, this lukewarm satisfaction is a red flag: **teams want to pay less, even if it means moving elsewhere.**

> **"I'd move [off Splunk] if I can, if it gets me to a lower cost point. I don't care where it's stored."**
>
> **Tony Hashem** | CISO & Director of IT Security, Vermont State Colleges System

In fact our data shows that **87% of leaders are either exploring or open to alternatives beyond their current Splunk infrastructure**.

Are you considering moving off Splunk due to cost concerns?

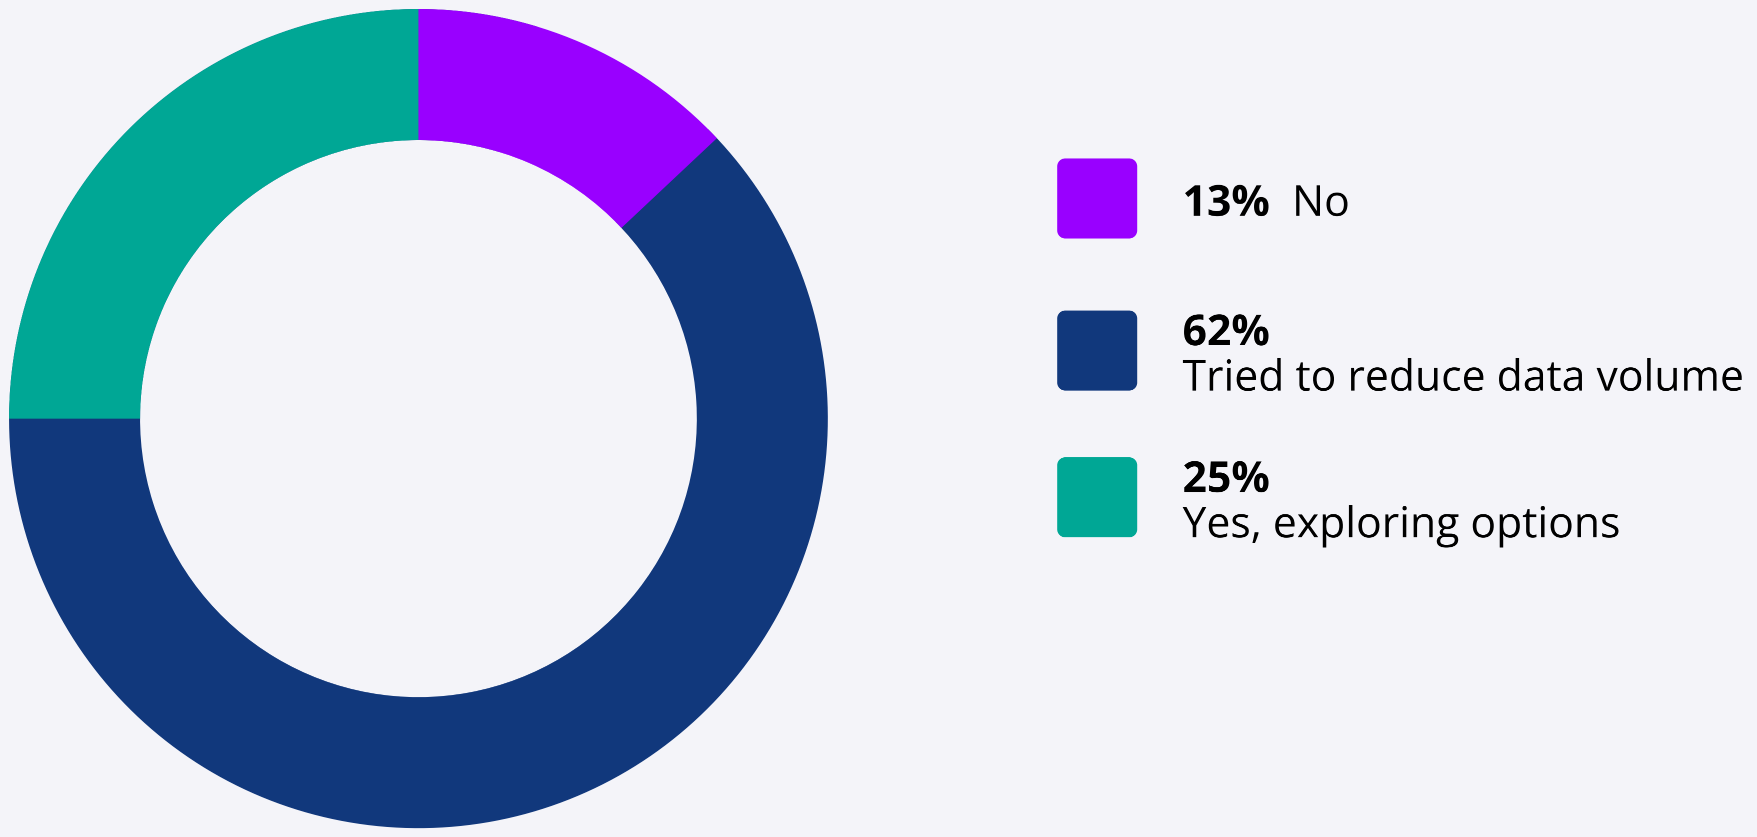

For now, most execs are reluctant to completely abandon Splunk and instead make difficult compromises to keep costs contained. But as data volumes rise, more leaders are preparing to act.

Practitioners were candid about this shift:

> **“Splunk is a powerful platform, but the economics don’t work at scale. We’ve started evaluating other approaches because we just can’t afford to keep throwing budget at the problem.”**

Chief Architect, Large Accounting and Tax Global Firm

> **“At this point, we’re not asking ‘if’ we’ll look for alternatives—it’s more about ‘when.’ The cost curve is unsustainable, and we need options.”**

Eike Schuster | Head of Cloud & Security Operations, Dentsply Sirona

Trends in executive behavior: Tough tradeoffs and seeking compatibility

Teams are forced to compromise

As costs spiral, nearly every executive is already implementing strategies to reduce their Splunk bill. These measures—whether filtering logs, archiving to cold storage, or limiting ingestion—save money in the short term but come at the expense of visibility, speed, and security.

> **“We can’t afford to keep everything. That means the visibility story gets weaker over time, which is the opposite of what observability should be.”**

Eike Schuster | Dentsply Sirona

In practice, most organizations are making **multiple compromises at once:** moving data to cheaper storage, dropping or filtering logs, and settling for summary indexes or pre-aggregated data. Instead of building the most effective observability systems, leaders are forced to choose whichever tactic does the least harm.

What strategies are you using today to control total cost of management?

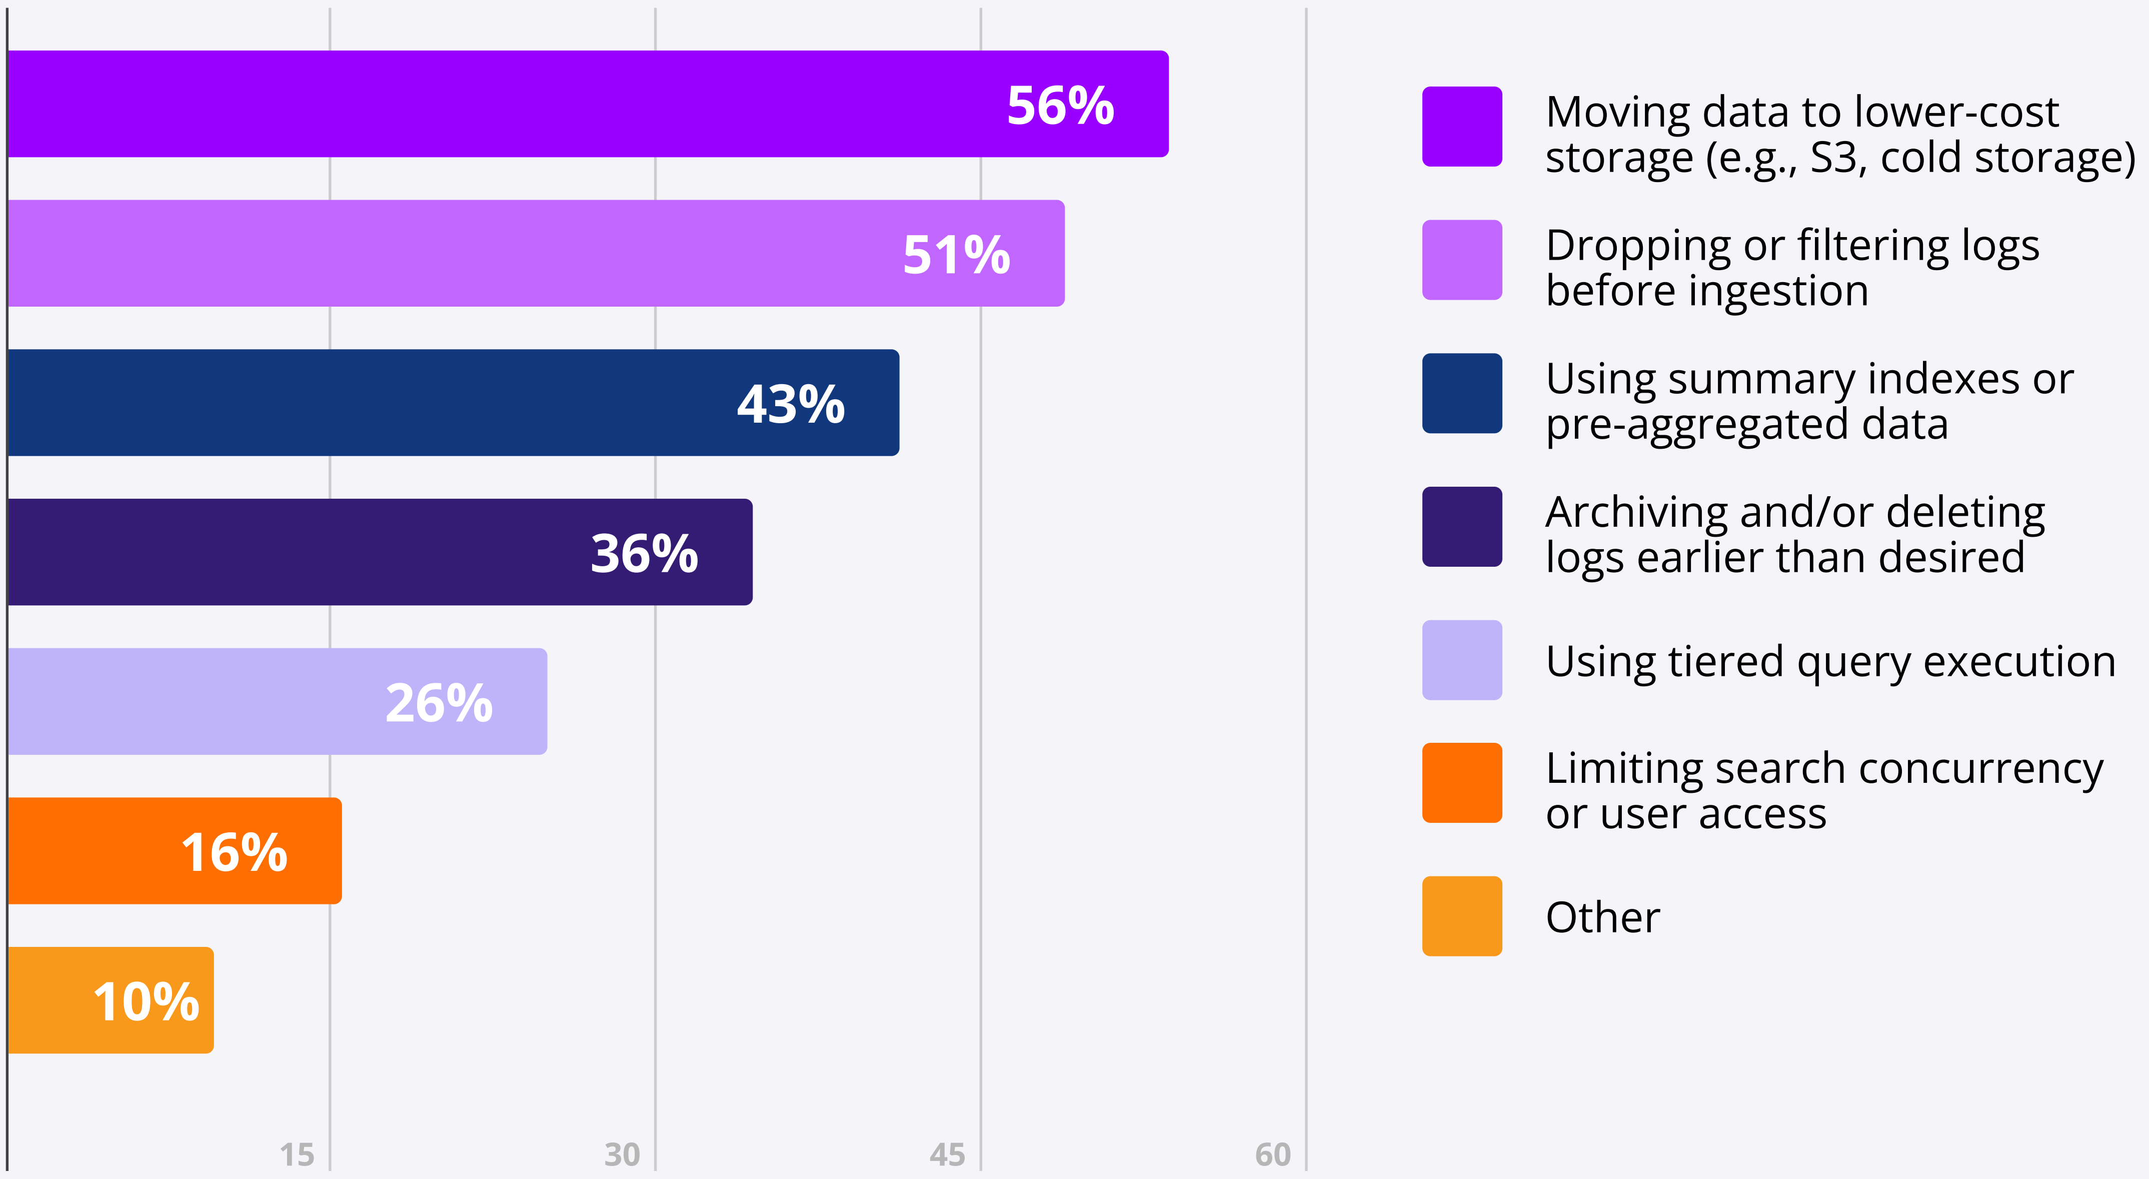

> **“The workarounds are endless: offloading to cheaper stores, limiting use cases, capping dashboards. It’s all about survival, not innovation.”**

Chief Architect, Large Accounting and Tax Global Firm

Stack complexity demands seamless integration

Splunk alone isn’t cutting it for most of those we surveyed, even with cost strategies in place.

92% of observability professionals use at least one other tool on top of Splunk.

The effort to cut Splunk costs by adopting new tools for specific use cases often backfires. It leads to fragmented, siloed data that's difficult to correlate, ultimately sacrificing data integrity.

> **“We use Splunk along with Graylog and Loki, but the problem is none of it is unified. It’s fragmented, hard to correlate, and visibility suffers. Teams end up spending more time stitching things together than solving problems.”**

Pooran Kumar | Director of IT & Digital Infrastructure, University College Birmingham

Compatibility matters

Observability leaders are in a bind: they don’t want to take on the risk and hassle of adding another solution, but they can’t make do with Splunk’s limitations. Their teams like Splunk’s interface and workflows—it’s familiar, trusted, and embedded in daily operations.

> **“Corporate growth is pushing us to get off Splunk [because of cost]. But our developers love Splunk. So that's the hustle going on right now. We have a runway of another six months to hope that they will move off Splunk”**

Head of IT Operations – Global Education and Publishing Company

That tension is driving openness to Splunk-compatible alternatives. Leaders want something that relieves the financial and operational pressure while letting teams continue to use the tools they know. When presented with a fully SPL-compatible option, 98% of leaders said they’d be open to adopting it.

> **“Splunk has value, but it’s just not sustainable at the scale we need. If there’s a way to keep the workflows while reducing cost, that’s exactly what we’re looking for.”**

Head of Cloud Site Reliability Engineering (Acting) – Global Investment Management Firm

From tradeoffs to transformation: Why we need a new model

The costs of observability today—both in terms of capital expenditure and employee time—are increasingly unsustainable: prices are rising, satisfaction is falling, and teams are stuck making uncomfortable compromises.

Cost optimization alone isn't enough to fix the root cause of these issues. Executives need a new approach that:

1. Breaks the link between cost and data volume.
2. Allows high-performance access to all data.
3. Works seamlessly with tools they already have.

It's time to decouple - The Imply Lumi solution

Imply Lumi: A modern observability layer

Imply Lumi is the industry's first **Observability Warehouse —** a high-performance data layer purpose-built for storing and querying observability data at scale.

By **decoupling storage and compute,** Imply Lumi breaks down silos and simplifies complex stacks—while giving you control over how and where data is stored, queried, and visualized.

Key benefits of Imply Lumi

01. Scale without runaway cost

Traditional systems force painful tradeoffs between cost and visibility. Imply Lumi is built for both **scale and efficiency.** With **schema-on-persist** and advanced compression, Imply Lumi shrinks your data footprint—delivering predictable costs while supporting full-fidelity retention, even at petabyte scale.

02. Fast queries on all your data

Today, many teams slow down investigations by shuffling data into cold storage just to keep costs in check. That means real-time queries turn into hours-long delays. Imply Lumi changes the game: its architecture treats **all data as fast data,** enabling interactive queries directly on compressed storage. The result is consistently faster queries than Splunk, whether for real-time troubleshooting, historical analysis, or cross-event correlation.

03. Compatibility without disruption

Teams don't want to rip and replace. That's why Imply Lumi is fully **Splunk-compatible:** it plugs directly into your existing SPL searches, dashboards, and alerts. You keep the workflows your teams know—while finally getting the scale, speed, and savings they've been asking for.

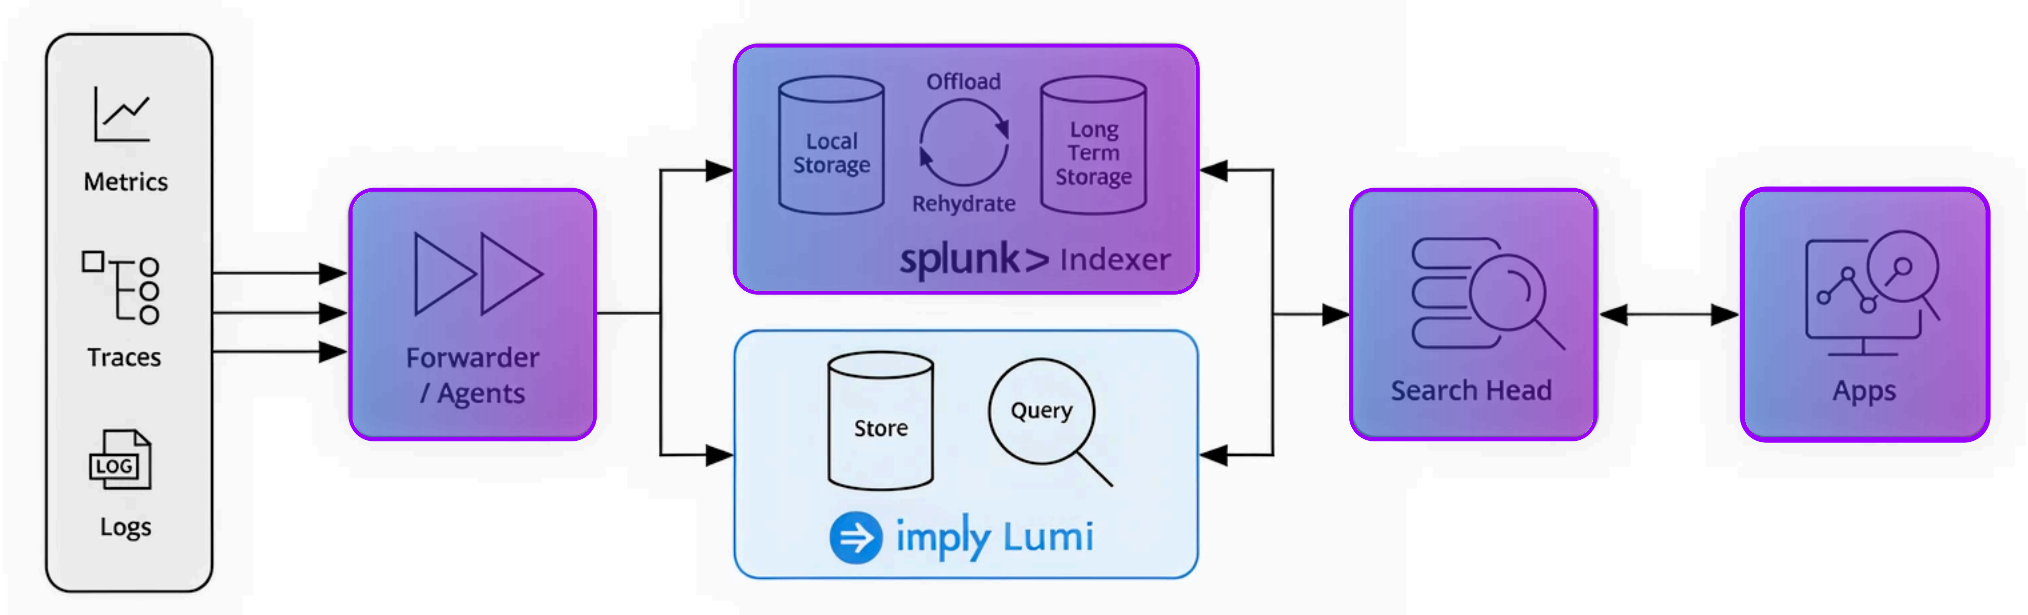

```
index = federated:index_name
```

"If something gave us Splunk-like search but with better economics, we'd be all over it. Compatibility is the only way change is realistic."

Head of Cloud Site Reliability Engineering (Acting) – Global Investment Management Firm

Conclusion

This report confirms what many leaders already recognize: observability costs are rising faster than value. Executives are being forced into trade-offs—sacrificing storage, performance, and visibility just to stay within budget. The result is a cycle of higher spend and diminishing returns.

Imply Lumi is built to break that cycle. Fully compatible with existing observability tools, Imply Lumi serves as a high-performance storage and compute layer that dramatically reduces cost while preserving the dashboards, agents, and workflows teams already rely on.

Why teams choose Imply Lumi:

Observability leaders face a difficult balance: they can't afford the disruption of adding new tools, but the limitations of legacy platforms are holding them back. Teams value the familiarity of their current workflows—what they need is a way to extend them.

Imply Lumi delivers that balance:

- Smaller storage footprint: Ingest and retain datasets that were previously cost-prohibitive, all at full fidelity
- Faster queries: Accelerate investigations across larger, more complex data without workflow changes
- Always-on data access: Analyze data continuously without filtering, archiving, or dropping events
- Zero Workflow Changes: Keep your existing agents, dashboards, and tools.

Contact us to learn how Imply Lumi can help you cut costs and unlock more from your observability Data.

About Imply

Imply, the Data Layer for Observability, Security, and AI, empowers organizations to keep more data, search it faster, and spend less — without changing their tools. Founded by the original creators of Apache Druid®, Imply delivers cloud-native infrastructure trusted by leading enterprises worldwide. Backed by a16z, Bessemer Venture Partners, and Thoma Bravo, Imply is headquartered in Silicon Valley with operations across North America, EMEA, and APAC. Learn more at imply.io.